Multi-Rover Coordination

Automated Reasoning PI Workshop
September 4, 2002

Reid Simmons, Stephen Smith,
Anthony Stentz, Jeff Schneider,
Dani Goldberg (presenter)

Robotics Institute
Carnegie Mellon University

Overview

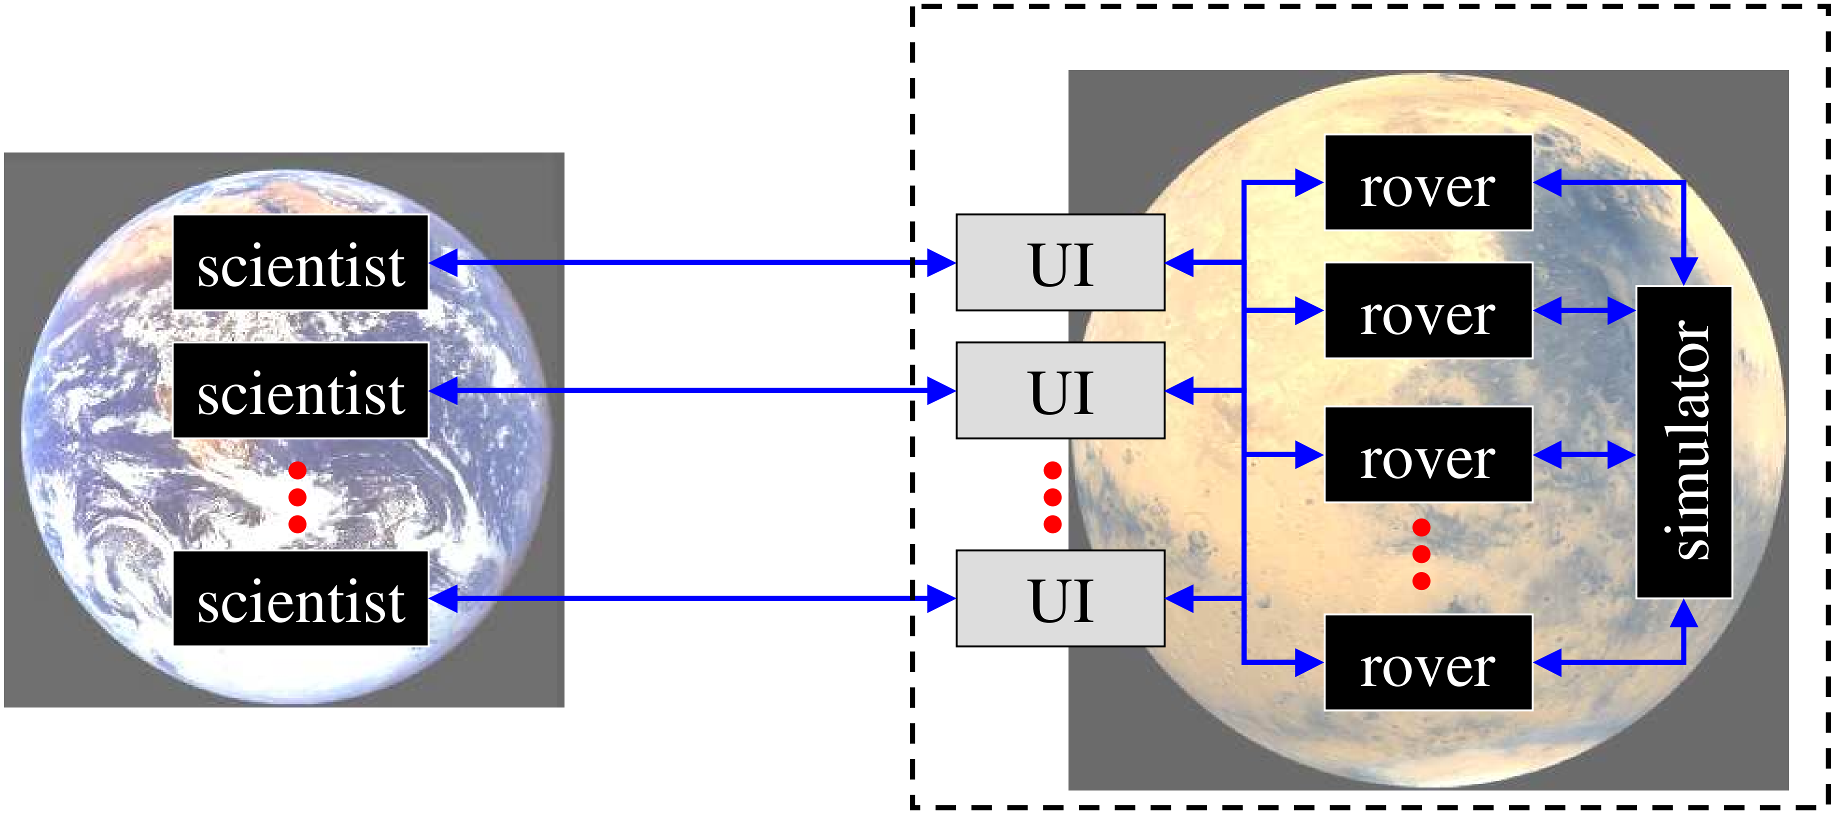

- many scientists
- many tasks/goals
- relative priorities

- bandwidth
- blackouts

- interface
- broker
- “OpTrader”

- heterogeneous
- autonomous
- limited resources

Scenario Overview

Year 1 Accomplishments

Distributed three-layered architecture

Multi-agent planning capability

Simulator

Scenarios

Distributed Three-Layered Architecture

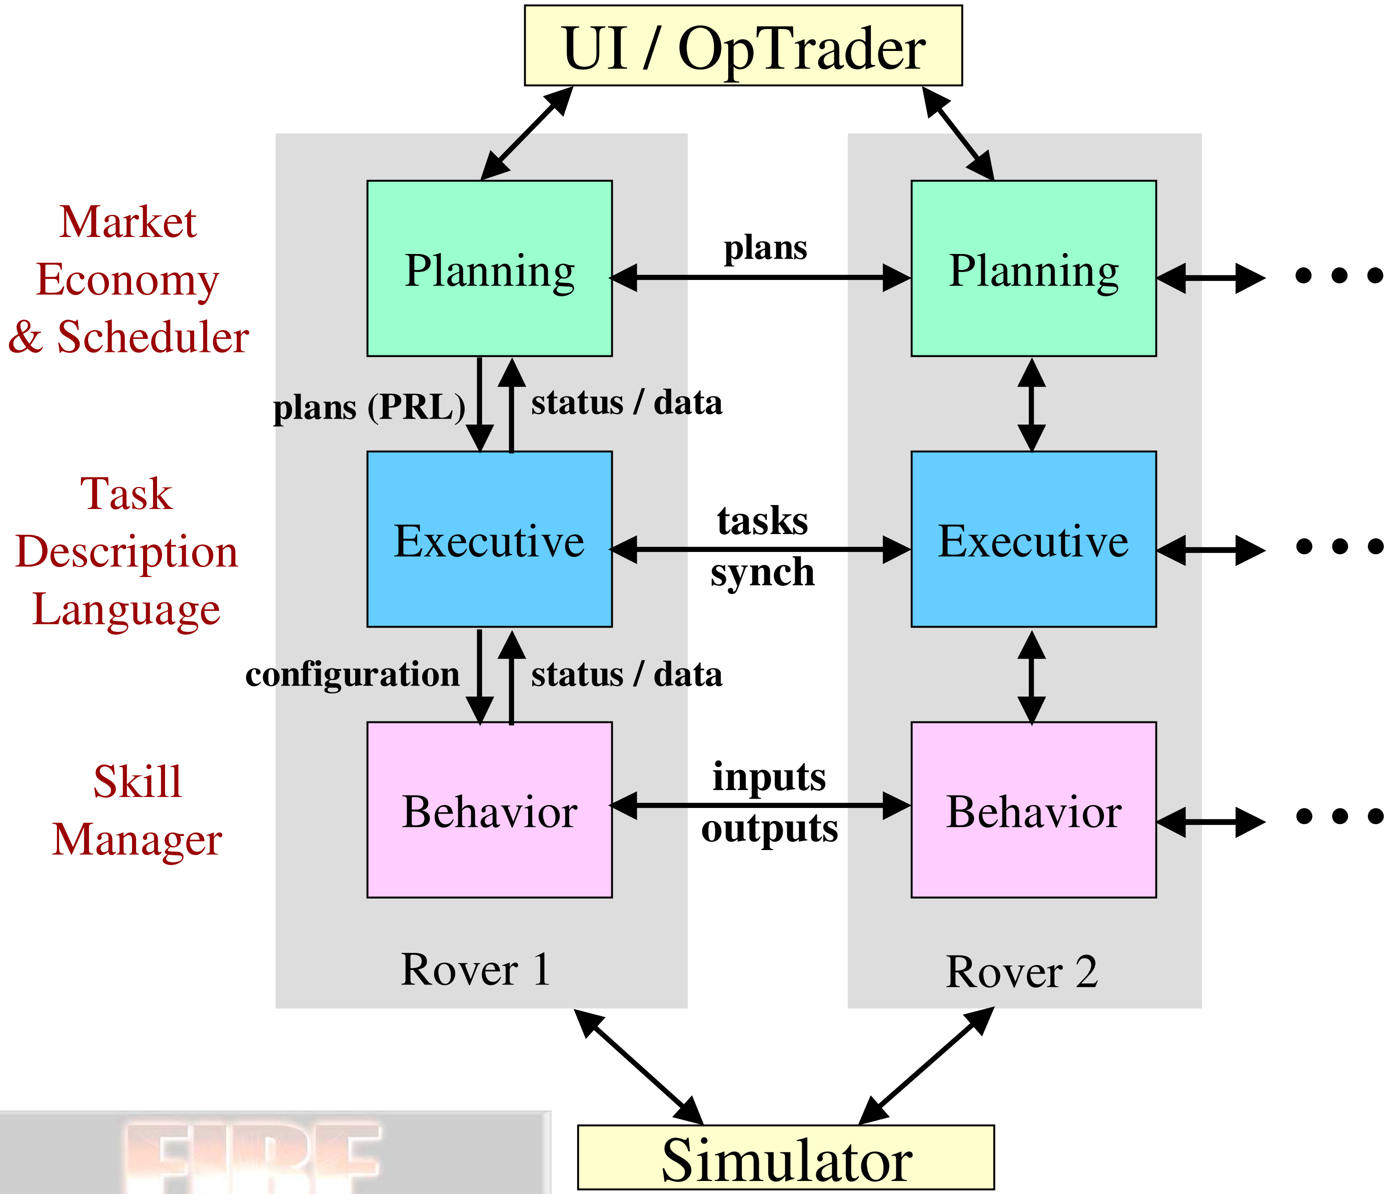

Control Flow
- Fixed data
- Fixed protocol

Data Flow (DDS)
- Define data
- Define protocol

Capability Server
- Propagation
- Reconciliation
- Configuration

Planned Use
- Reconfiguration
- Failures
- Distributed tasks

Distributed Behavioral and Executive Layers

Behavioral Layer

- Real-time sensor/effector feedback loops
- Efficient distributed servo loops (sensor/effector connections)
- Tightly coordinated controllers (effector/effector connections)

Executive Layer

- Mediator between planning and behavioral layers
  - Receives distributed plans, translates them, establishes appropriate behavioral connections/functionality
- Task Description Language (TDL)
  - Hierarchical task decomposition, task synchronization, execution monitoring, exception handling
  - Recent improvements transparently handle task-level coordination between robots using message passing

Future Extensions: full support of distributed tasks

Multi-Agent Planning Capability

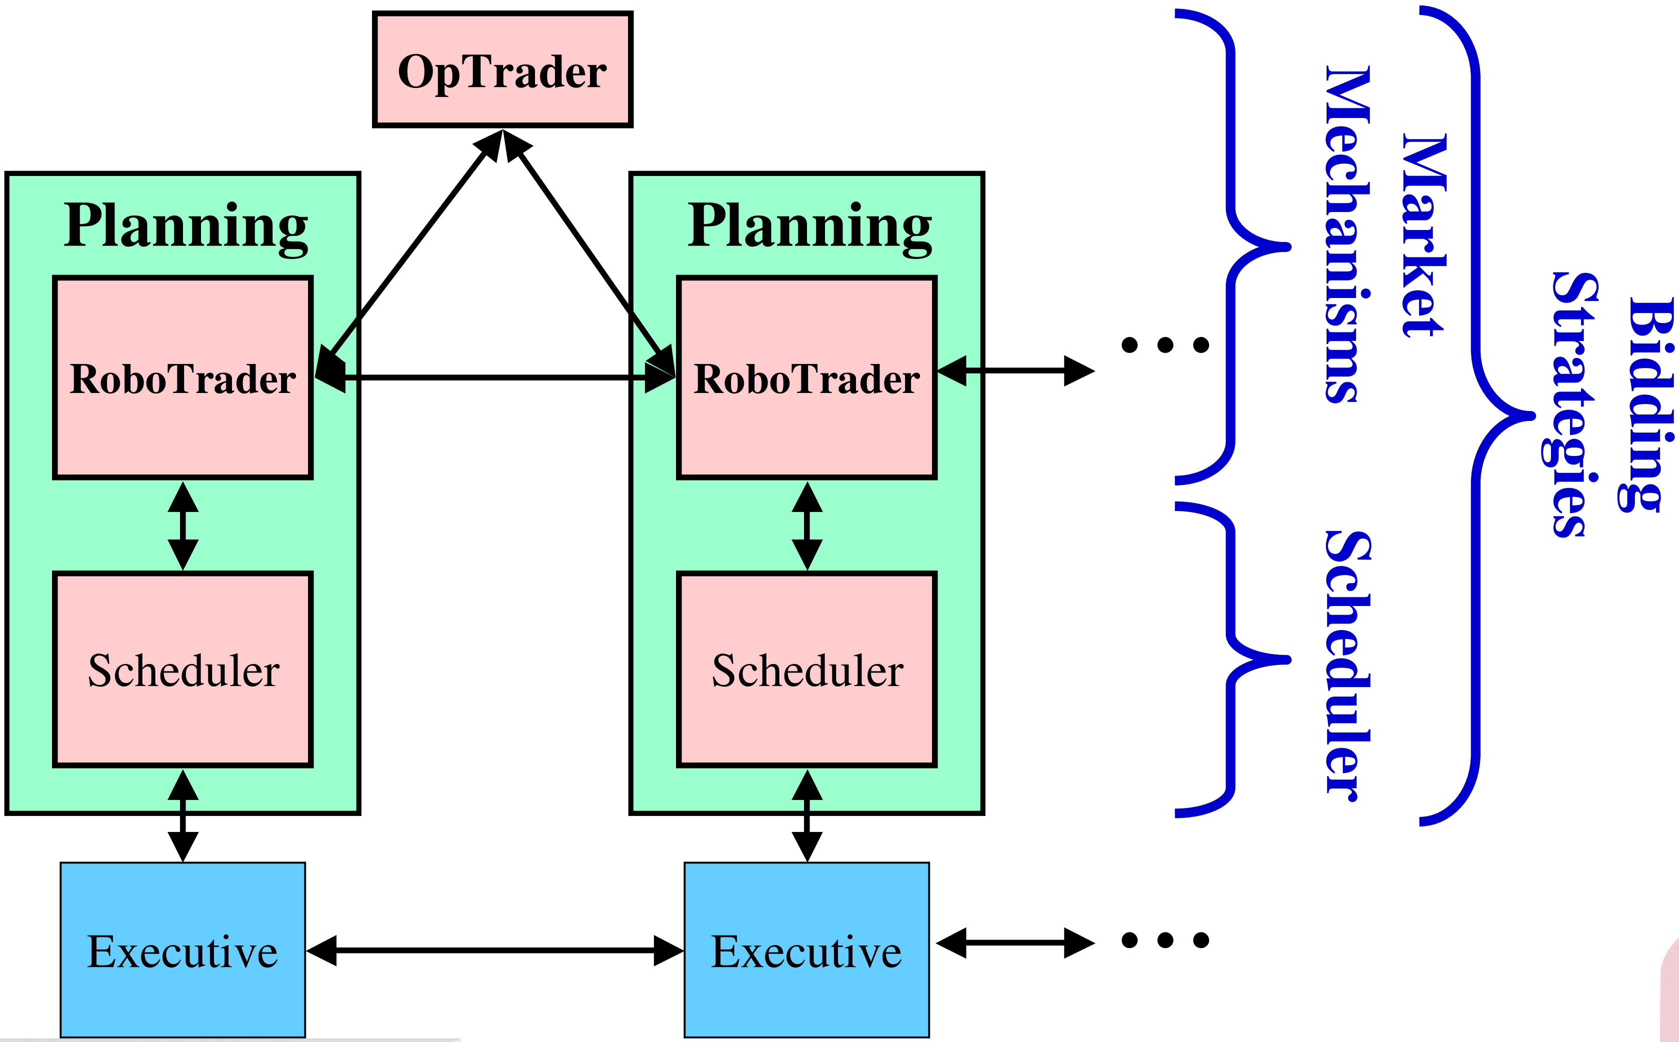

Market Mechanisms

- Inter-Trader bidding
  - One Trader (auctioneer) make a call for bids on a task/tasks
  - Other Traders (bidders) send bids with costs
  - Auctioneer awards task(s) to bidder(s) with lowest costs
- Traders are continuously auctioning and bidding tasks
  - Tasks that are up for bid should not be executed
  - A Trader can simultaneously be a bidder and auctioneer only if the tasks it is bidding on are awarded before its own auction terminates

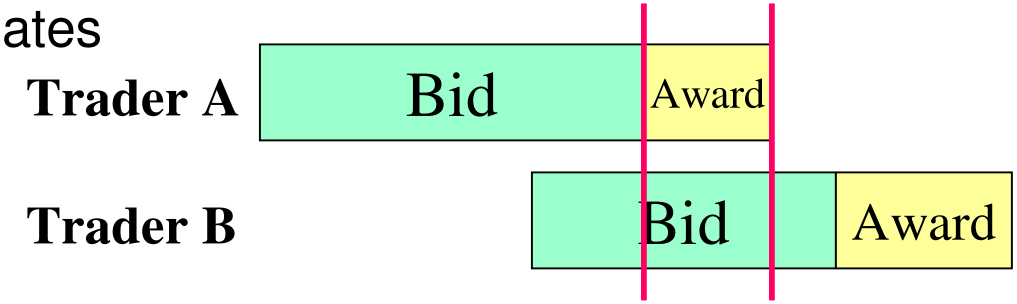

Market Mechanisms: Leaders for Opportunistic Optimization

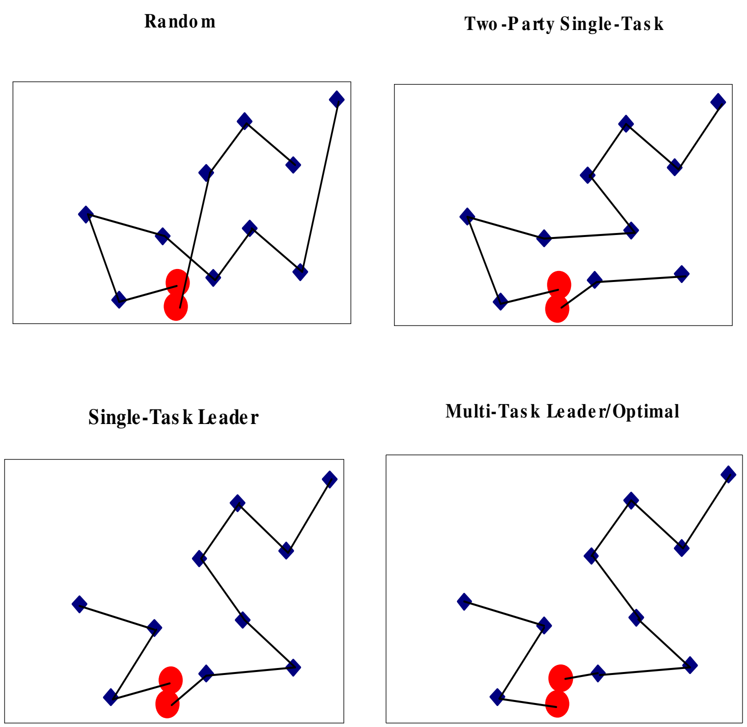

Distributed TSP Problem

- Leaders generate plans requiring multi-agent participation and “sell” them in the market place.
- Cost-effective plans win in the market since they generate enough profit to “buy” participation.
- The optimization is opportunistic and local to the participating agents.
- The first leader implemented is a combinatorial exchange for multi-party and multi-task deals.

Market Mechanisms: Task Clustering

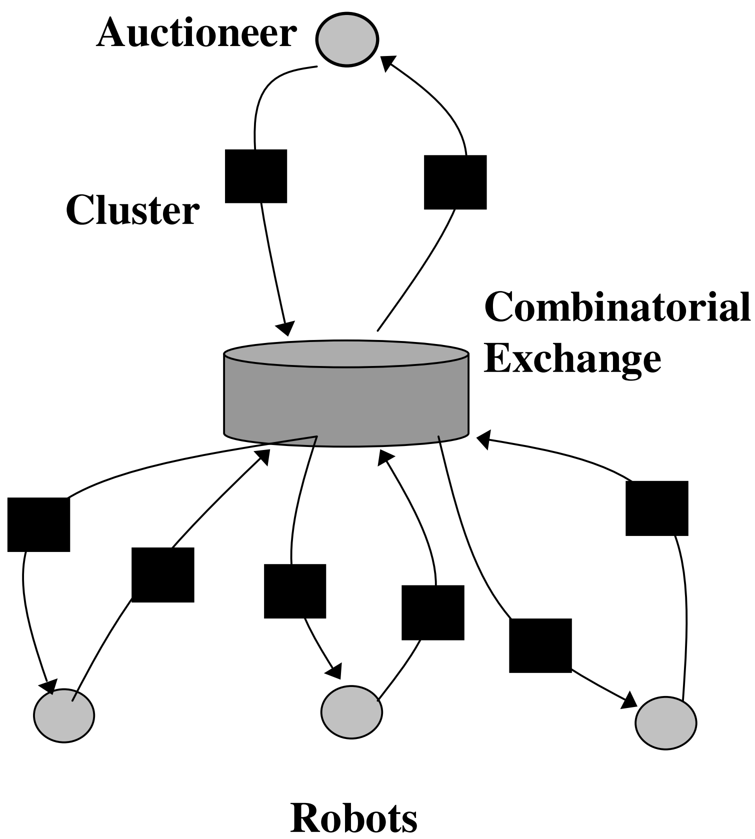

- Tasks which can be executed cost effectively as a group are clustered.
- Agents form clusters and “buy/sell” them on the market.
- Clusters reduce the combinatorics of matching tasks to agents.

Constructing Bids and Managing Execution

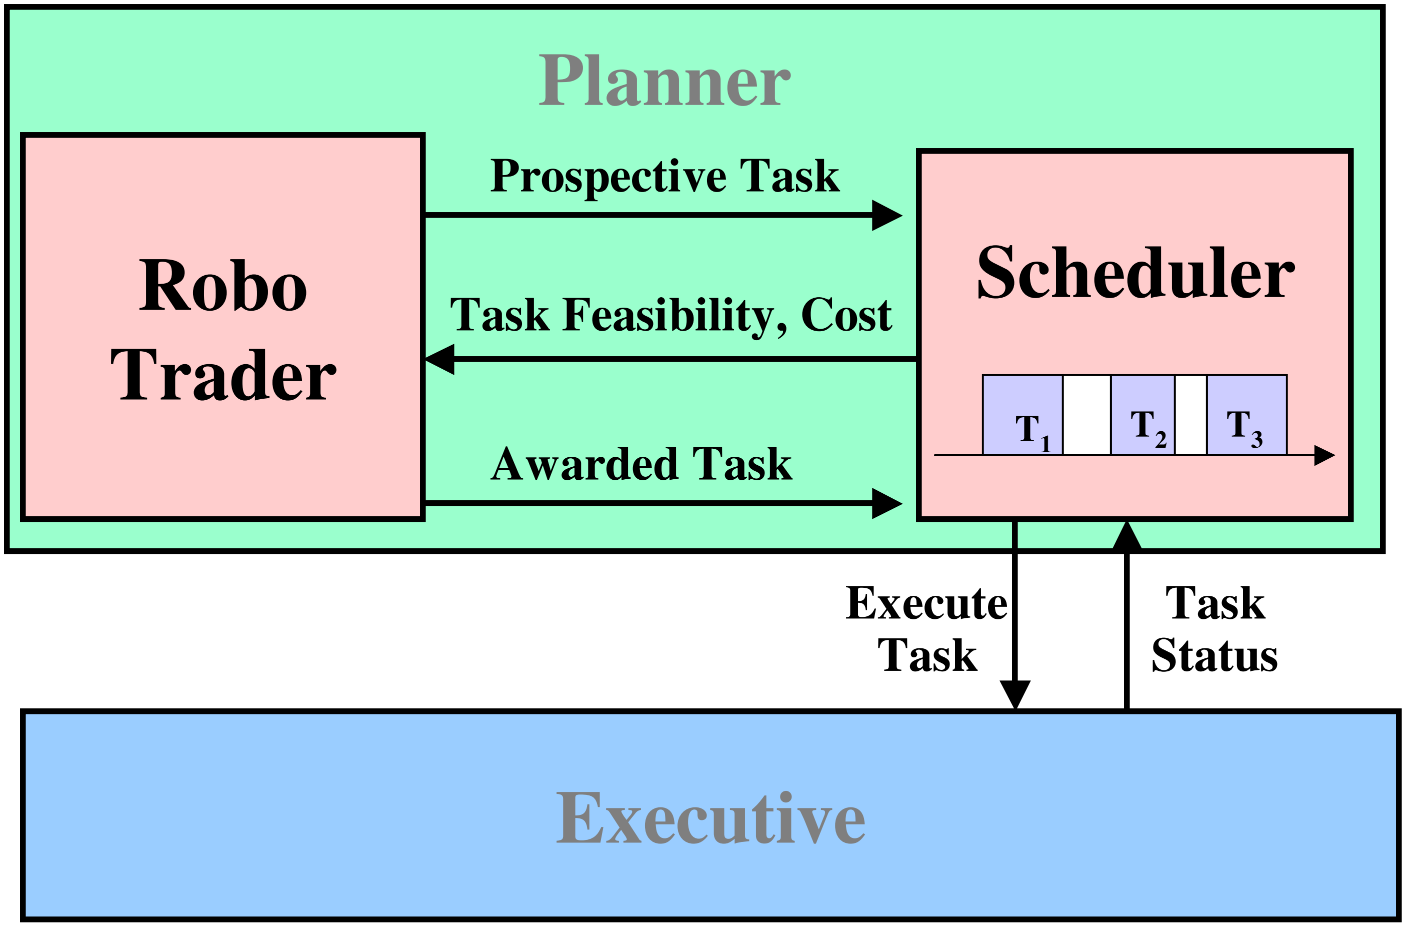

Scheduler: Stochastic Amplification of Scheduling Search Heuristics

Basic Idea (Cicirello & Smith, CP-02)

- Assume good scheduling heuristic
- Introduce randomness in the search when heuristic is less discriminating

Desirable Characteristics

- Good solutions fast, with anytime properties if time permits
- Flexibility to incorporate a range of different heuristics

Current Assumption: Scheduling problem formulated as TSP

Planned Extensions

- Deadlines; time-windows
- Coordinated multi-rover tasks
- More sophisticated cost drivers
- Additional resource constraints

Current Bidding Strategy: Estimated Time Cost

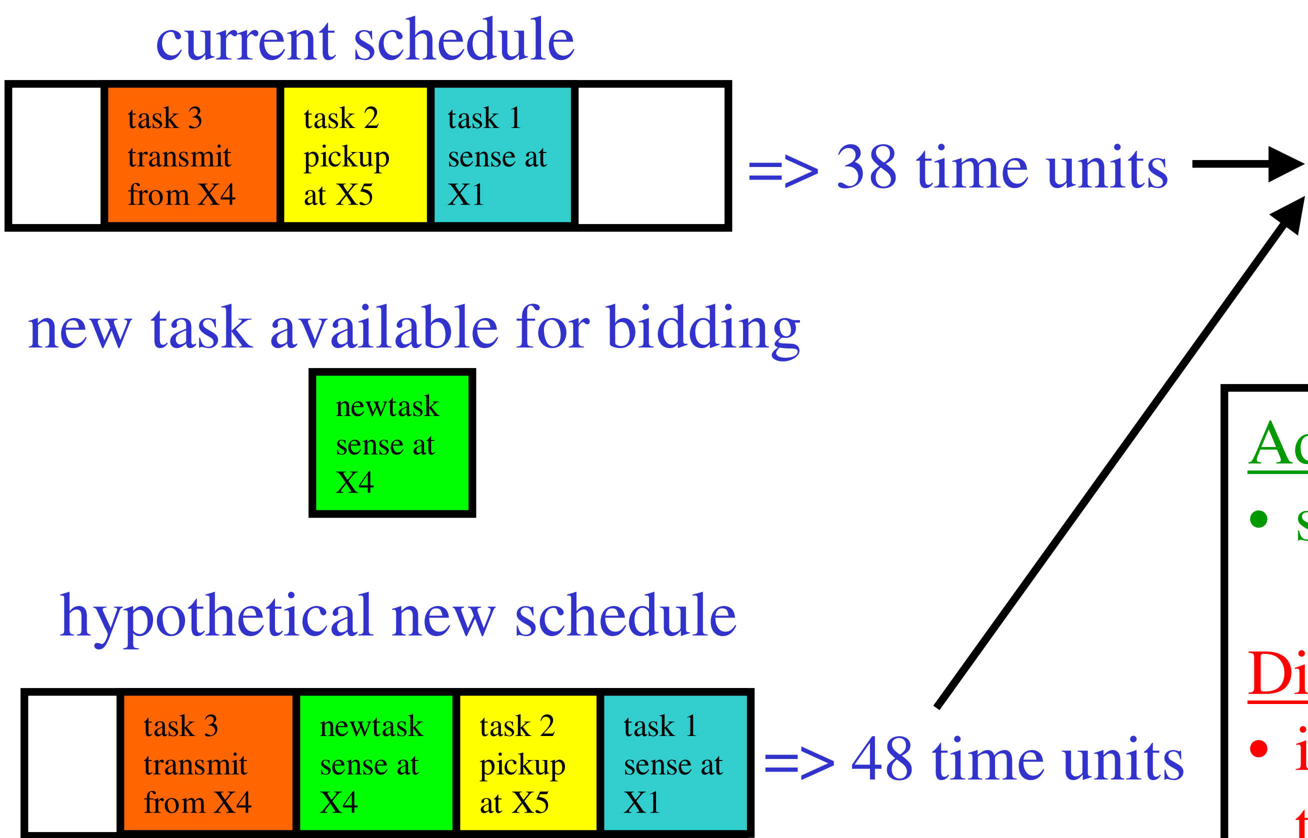

BID 10

Task is awarded to lowest bidder

Advantages:
- simple to compute

Disadvantages:
- ignores the value of the task
- ignores the value of other tasks that may be on offer

Future Bidding Strategy: Estimated Opportunity Cost

Learn an "expected revenue" function approximator with data from simulation runs

features of current schedule → **Function Approximator** → expected total revenue

features of available tasks → **Function Approximator**

current schedule

	Task3: 8 transmit from X4	Task2: 5 pickup at X5	Task1: 6 sense at X1	

=> estimated revenue: 24 → BID 24 -19 = 5

lowest bid is accepted if it does not exceed task value

new task available for bidding

newtask sense at X4

other remaining tasks

sense at X8 value 9

pickup at X3 value 8

value 7

X8 value 9

hypothetical new schedule (do new task for free)

	Task3: 8 transmit from X4	Task: 0 sense at X4	Task2: 5 pickup at X5	Task1: 6 sense at X1

=> estimated revenue: 19

Remaining issues:

- choice of function approximator
- efficient computation

In Action…

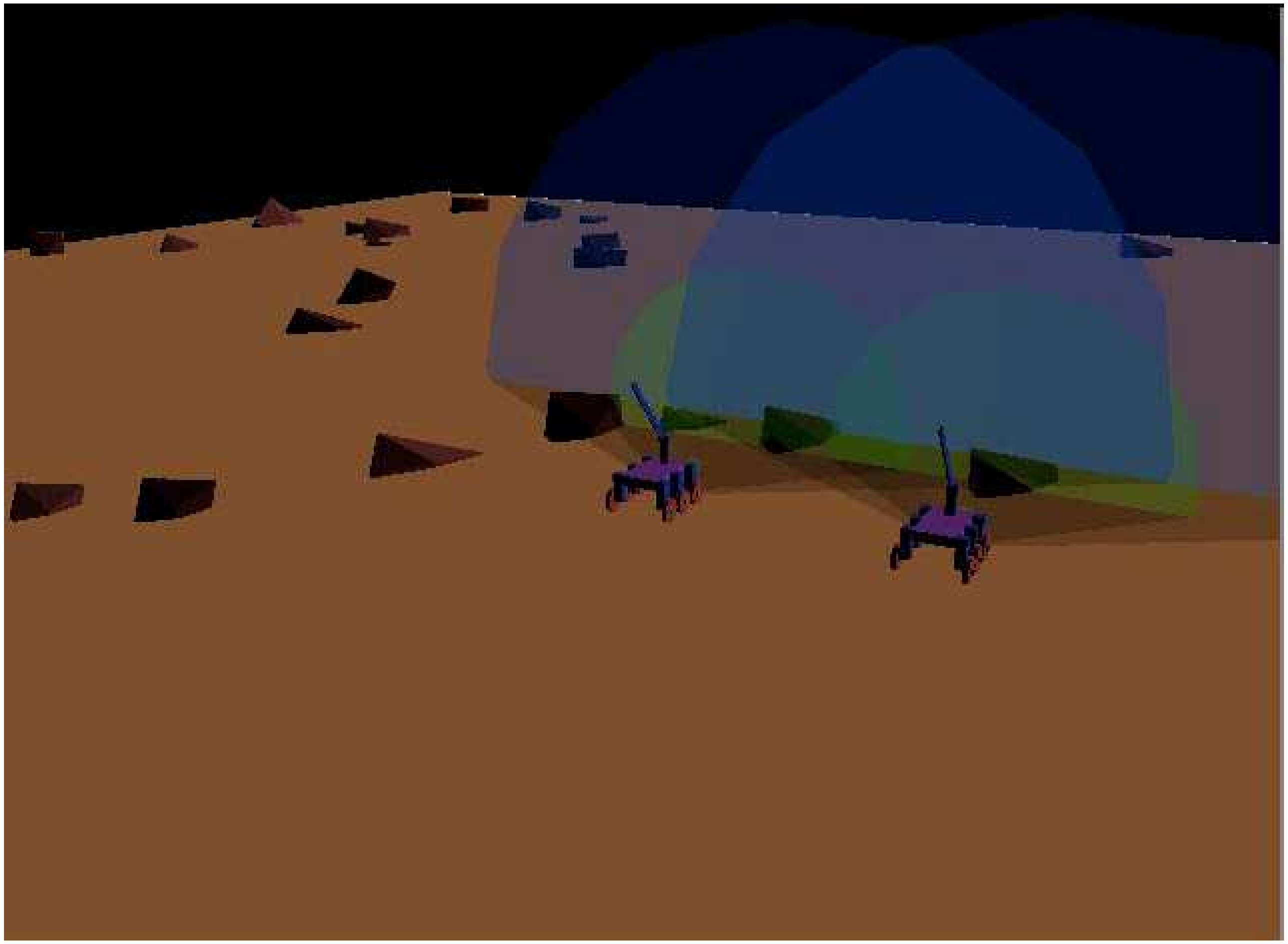

Continuing Work

- "ComTask" problem
  - Maintain a communications network between the rovers based on limited range, line of sight communications equipment.
- Planning based on the value of information
  - Plan to maximize information gain
- Experimental Plan
  - Experiments
    - Vary market parameters
    - Vary task structure
    - Vary environment
  - Performance Data
    - % rocks collected, activity of agents, total cost decline

Personnel

PI's

- Reid Simmons
- Anthony Stentz
- Stephen Smith
- Jeff Schneider

Postdocs

- Dani Goldberg

Research Programmers

- David Apfelbaum

Graduate Students

- Drew Bagnell
- Vincent Cicirello
- M. Bernardine Dias
- Maayan Roth
- Brennan Sellner
- Trey Smith

Undergraduate Students

- Stuart Anderson

www.cs.cmu.edu/~fire

Heterogeneous Multi-Rover Coordination

Reid Simmons, Steve Smith, Tony Stentz, Jeff Schneider
Carnegie Mellon University

Goal: Provide capabilities for multiple, heterogeneous robots to coordinate activities autonomously, efficiently, and reliably.

Objectives:

- Develop architectural framework for exploring tradeoffs between planning, negotiation, sensing, and action.
- Demonstrate utility of heterogeneous teams of robots to perform coordinated tasks autonomously and reliably.

Key Innovations:

- Market-based architecture for team formation, resource allocation.
- Distributed global planning and task execution.
- Explicit tradeoffs between planning, sensing, and execution.
- Probabilistic reasoning and learning algorithms for sensing and task execution.

NASA Relevance:

- Future space missions, such as Mars exploration or in-space assembly, will benefit from teams of autonomous robots. Heterogeneous teams will be more capable and reliable than similarly sized homogeneous robots.

Accomplishments to date:

- Developed multi-rover graphical simulator.
- Implemented initial, coordinated task execution functionality.
- Defined/implementing market-based architectural framework.
- Provided initial scheduling/planning capabilities.
- Developed detailed Mars rover exploration scenario.
- Demonstrated market/scheduler interaction for task allocation.

Description

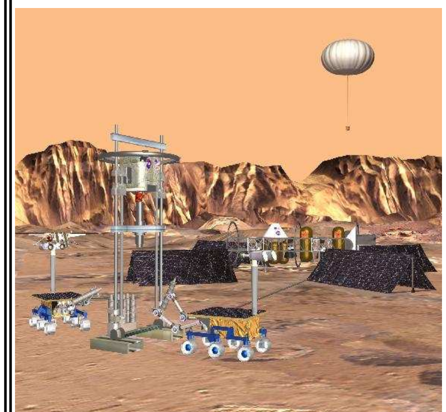

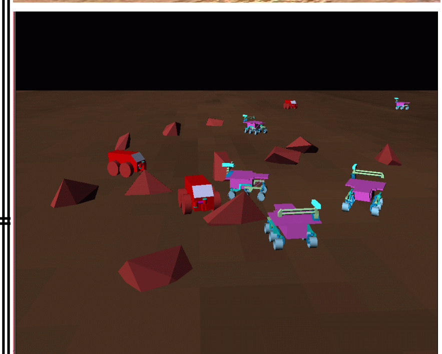

- **Market-Based Architecture**: Use distributed bidding techniques to form teams, allocate tasks, and deal with changing conditions.
- **Global Coordination**: Use both distributed and centralized planning and scheduling to efficiently create highly optimal plans for multiple agents.
- **Coordinated Task Execution**: Use distributed task synchronization and monitoring to enable teams of robots to perform tasks that none can do individually.
- **Valuing Sensor Actions and Plans**: Use information and decision theoretic techniques to make trade offs in when, what, and how to sense, and cost/value potential plans.

Schedule:

- FY 02: Distributed multi-layer architecture; multi-rover planning capability.
- FY 03: Extension to complex tasks requiring explicit multi-rover coordination in planning and execution.
- FY 04: Incorporate value of information into planning; expansion to more complex scenarios; experimentation and validation.